arXiv:1207.4959v1 [stat.ME] 20 Jul 2012

Aggregating density estimators: an empirical study

M. Bourel[*], B. Ghattas[†]

July 23, 2012

Abstract

We present some new density estimation algorithms obtained by bootstrap aggregation like Bagging. Our algorithms are analyzed and empirically compared to other methods found in the statistical literature, like stacking and boosting for density estimation. We show by extensive simulations that ensemble learning are effective for density estimation like for classification. Although our algorithms do not always outperform other methods, some of them are as simple as bagging, more intuitive and has computational lower cost.

Keywords: Machine Learning, Histogram, Kernel Density Estimator, Bootstrap, Bagging, Boosting, Stacking.

1 Introduction

Ensemble learning is one of the most challenging recent approaches in statistical learning. Bagging ([1]), Boosting ([5]), Stacking ([2]), and Random forests ([3]) have been declared to be the best of the chelf classifiers achieving very high performances when tested over tens of various datasets selected from the machine learning benchmark. All these algorithms had been designed for supervised learning, sometimes initially restricted to regression or binary classification. Several extensions are still under study: multivariate regression, multiclass learning, and adaptation to functional data or time series.
Very few developments exist for ensemble learning for unsupervised framework, clustering analysis and density estimation. Our work concerns the latter case which may be seen as a fundamental problem in statistics. Among the last developments, we found some extensions of boosting ([5]) and stacking ([11]) to density estimation.

[*]M. Bourel, mbourel@fing.edu.uy, IMERL, Facultad de Ingenieria, Universidad de la República, Julio Herrera y Reissig 565, 11300 Montevideo, Uruguay.

[†]B. Ghattas, badih.ghattas@univ-amu.fr, Université d'Aix Marseille, Institut de Mathématiques de Luminy, case901, 163 avenue de Luminy, marseille 13009, France

In this paper we suggest some simple algorithms for density estimation in the same spirit of bagging and stacking where the weak learners are histograms. We show by extensive simulations that aggregation gives rise to effective better estimates. We compare our algorithms to several algorithms for density estimation, some of them are simple like Histogram and kernel density estimators (KDE) and others rather complex like stacking and boosting which will be described in details. As we will show in the experiments we do, although the accuracy of our algorithm is not systematically higher than other ensemble methods, it is with no doubt simpler, more intuitive and computationally less expensive. Boosting algorithms and stacking for density estimation are described in section 2. Section 3 describes our algorithms. Simulations and results are given in section 4 and concluding remarks and future work in section 5.

2 A review of the existing algorithms

In this section we review some density estimators obtained by aggregation. They may be classified in two categories depending on the aggregation form.
The first type has the form of linear or convex combination:

$$f_M(x) = \sum_{m=1}^{M} \alpha_m g_m(x) \tag{1}$$

where g_m is typically a parametric or non parametric density model, and in general different values of m refer typically to

- different parameters values in the parametric case or,
- different kernels, or
- different bandwidths for a chosen kernel for the kernel density estimators.

The second type of aggregation is multiplicative and is based on the ideas of high order bias reduction for kernel density estimation ([6]). The aggregated density estimator has the form:

$$f_M(x) = \prod_{m=1}^{M} \alpha_m g_m(x) \tag{2}$$

2.1 Linear or convex combination of density estimators

This kind of estimators (1) has been used in several works with different construction schemes.

- In [10], [8] and [12] the weak learners g_m are introduced sequentially in the combination. At step m, g_m is chosen to maximize the log likelihood of

$$f_m(x) = (1-\alpha) f_{m-1}(x) + \alpha g_m(x) \tag{3}$$

where g_m is a density selected among a fixed class $\mathcal{H}$.
In [10] g_m is selected among a non parametric family of estimators, and in [8] and [12], it is taken to be a Gaussian density or a mixture of Gaussian densities whose parameters are estimated. Different methods are used to estimate both density g_m and the mixture coefficient α. In [8] g_m is a Gaussian density and the log likelihood of (3) is maximized using a special version of Expectation Maximization (EM) taking into account that a part of the mixture is known.

The main idea underlying the algorithms given by [10] and [12] is to use Taylor expansion around the negative log likelihood that we wish to minimize:

$$\sum_i -\log(f_m(x_i)) = \sum_i -\log(f_{m-1}(x_i)) - \alpha \sum_i \frac{g_m(x_i)}{f_{m-1}(x_i)} + O(\alpha^2)$$

For α small we have the approximation

$$\sum_i -\log(f_m(x_i)) \sim \sum_i -\log(f_{m-1}(x_i)) - \alpha \sum_i \frac{g_m(x_i)}{f_{m-1}(x_i)}$$

thus, minimizing the left side term is equivalent to maximizing $\sum_i \frac{g_m(x_i)}{f_{m-1}(x_i)}$.

All the algorithms described above are sequential and the number of weak learners aggregated may be fixed by the user.

- [11] use stacked density estimator applying the same aggregation scheme as in stacked regression and classification ([13]). The M densities g_m are fixed in advance (KDE with different bandwidths). The data set $\mathcal{L} = \{x_1, \ldots, x_n\}$ is divided into V cross validation subsets $\mathcal{L}_1, \ldots, \mathcal{L}_V$. For $v = 1, .., V$, denote $\mathcal{L}^{(-v)} = \mathcal{L} - \mathcal{L}_v$. The M models $g_1, \ldots, g_M$ are fitted using the training samples $\mathcal{L}^{(-1)}, \ldots, \mathcal{L}^{(-V)}$, the obtained estimates are denoted $\widehat{g}_m^{(-1)}, \ldots, \widehat{g}_m^{(-V)}$ for all $m = 1, \ldots, M$. These models are then evaluated over the test sets $\mathcal{L}_1, \ldots, \mathcal{L}_V$, getting the vectors $\widehat{g}_m^{(-v)}(\mathcal{L}_v)$ for $m = 1, \ldots, M$, $v = 1, \ldots, V$ put within a $n \times M$ matrix

$$A = \begin{pmatrix} \widehat{g}_1^{(-1)}(\mathcal{L}_1) & \ldots & \ldots & \widehat{g}_M^{(-1)}(\mathcal{L}_1) \\ \widehat{g}_1^{(-2)}(\mathcal{L}_2) & \ldots & \ldots & \widehat{g}_M^{(-2)}(\mathcal{L}_2) \\ \vdots & & & \vdots \\ \widehat{g}_1^{(-V)}(\mathcal{L}_V) & \ldots & \ldots & \widehat{g}_M^{(-V)}(\mathcal{L}_V) \end{pmatrix}$$

This matrix is used to compute the coefficients $\alpha_1, \ldots, \alpha_M$ using the Expectation-Maximization algorithm. Finally, for the output model, we re-estimate the individual densities $g_1, \ldots, g_M$ from the whole data $\mathcal{L}$.

- In [9] the densities $\{g_m\}_{m=1,\ldots,M}$ are fixed in advance like for stacking (KDE estimators with different bandwidths). The dataset is split in two parts. The first sample is used to estimate the densities g_m, whereas the coefficients α_m are optimized using the second sample. The splitting process is repeated and the aggregated estimators for each data split are averaged. The final model has the form

$$f_M(x) = \frac{1}{card\{S\}} \sum_{s \in S} \tilde{g}_M^s(x)$$

where S is the set of all the splits used and

$$\tilde{g}_M^s(x) = \sum_{m=1}^{M} \alpha_m g_m^s(x)$$

is the aggregated estimator obtained from one split s of the data, g_m^s is the individual kernel density function estimated over the learning sample obtained from the split s. This algorithm is called *AggPure*.

2.2 Multiplicative aggregation

The only algorithm giving rise to this form of aggregation is the one described in [4] called *BoostKde*. It is a sequential algorith where at each step m the weak learner is computed as follows:

$$\hat{g}_m(x) = \sum_{i=1}^{n} \frac{w_m(i)}{h} K\left(\frac{x - x_i}{h}\right)$$

where K is a fixed kernel, h its bandwidth, and $w_m(i)$ the weight of observation i at step m. Like for boosting, the weight of each observation is updated:

$$w_{m+1}(i) = w_m(i) + \log\left(\frac{\hat{g}_m(x_i)}{\hat{g}_m^{(-i)}(x_i)}\right)$$

where $\hat{g}_m^{(-i)}(x_i) = \sum_{j=1, j\neq i}^{n} \frac{w_m(j)}{h} K\left(\frac{x_j - x_i}{h}\right)$. The output is given by $\hat{f}_M(x) = C \prod_{m=1}^{M} \hat{g}_m(x)$, where C is a normalization constant. The Algorithm is resumed in figure 1.

1. For $i = 1, \ldots, n$, initialize the weigths of the observations $w_1(i) = \frac{1}{n}$ and fix the bandwidth h.

2. For $m = 1$ to M:

 (a) Compute the weighted kernel estimate

 $$\widehat{g}_m(x) = \sum_{i=1}^{n} \frac{w_m(i)}{h} K\left(\frac{x - x_i}{h}\right)$$

 (b) Update the weights $w_{m+1}(i) = w_m(i) + \log\left(\frac{\widehat{g}_m(x_i)}{\widehat{g}_m^{(-i)}(x_i)}\right)$

 with $\widehat{g}_m^{(-i)}(x_i) = \sum_{j=1, j\neq i}^{n} \frac{w_m(j)}{h} K\left(\frac{x_j - x_i}{h}\right)$.

3. Output: $\hat{f}_M(x) = C \prod_{m=1}^{M} \widehat{g}_m(x)$ (C normalization constant such that f_M integrates to unity).

Figure 1: Boosting kernel density estimation algorithm ((BoostKde), [4])

3 Aggregating Histograms

We suggest three new density estimators obtained by linear combination like in (1), all of them use histograms as weak learners. The first two algorithms may be parallelized and randomize the histograms. The third one is just a modification of Stacking.

The first algorithm is similar to Bagging ([1]). At each step m, a bootstrap sample of the original dataset is generated and g_m is the histogram obtained from the generated bootstrap sample with a fixed number of equally spaced breakpoints. We will refer to this algorithm as *BagHist*, it is detailed in figure 2.

The second algorithm (*AggregHist*) works as follows: let g_0 be the histogram obtained with the data set at hand using equally spaced breakpoints denoted $\mathcal{B} = b_1, b_2, ..., b_L$. Each weak learner g_m is an histogram constructed over the same initial data set but using a randomly modified set of breakpoints; γ is a tuning parameter which controls the variance of the perturbations. The algorithm is detailed in figure 3.

Finally, we introduce a third algorithm called *StackHist* where we replace in the stacking algorithm described in the previous section, the six kernel density estimators by three histograms with fixed number of breaks.

The values of the parameters used in these algorithms will be optimized. The procedure used will be described in the experiments section.

1. Let $\mathcal{L}$ be the original sample
2. For $m = 1$ to M:
 (a) Let $\mathcal{L}^m$ be a bootstrap sample of $\mathcal{L}$
 (b) Set g_m to be the histogram constructed over $\mathcal{L}^m$ with equispaced L breakpoints.
3. Output: $f_M(x) = \frac{1}{M} \sum_1^M g_m(x)$

Figure 2: Bagging of histograms (*BagHist*)

1. Let $\mathcal{L}$ be the original sample, g_0 be the histogram constructed over $\mathcal{L}$ and $\mathcal{B} = \{b_1, b_2, ..., b_L\}$ the set of the ordered optimized breakpoints.
2. For $m = 1$ to M:
 (a) Set $\mathcal{B}^m = \{b^*_{(1)}, b^*_{(2)}, ..., b^*_{(L)}\}$ the modified breakpoints obtained by setting $b^*_l = b_l + \varepsilon_l$ where $\varepsilon_l \sim N(0, \sigma)$ and $\sigma = \gamma \; min_{1<l\leq L} \{b_l - b_{l-1}\}$.
 (b) Set g_m to be the histogram constructed over $\mathcal{L}$ using the breakpoints $\mathcal{B}^m$.
3. Output: $f_M(x) = \frac{1}{M} \sum_1^M g_m(x)$

Figure 3: Aggregating histograms based on randomly perturbed breakpoints (*AggregHist*)

4 Experiments

We test several simulation models based on classical distributions and mixture models mostly used in the cited works and algorithms described above. The sample size is fixed at $n = 100, 500, 1000$.
We first show the estimates obtained using *BagHist* and *AggregHist* and analyze the effect of the number M of histograms aggregated, and then compare them to the other algorithms.

Up to our knowledge the existing algorithms for density estimation by aggregation have never been compared over a common benchmark simulation data.

4.1 Models used for the simulations

We denote by $\mathcal{M}1, \ldots, \mathcal{M}11$ the different simulation models which will be grouped according to their difficulty level.

- Some standard densities used in [4]:

 ($\mathcal{M}1$): standard Gaussian density $N(0,1)$

 ($\mathcal{M}2$): standard exponential density $f(x) = \begin{cases} 0 & x < 0 \\ e^{-x} & x \geq 0 \end{cases}$

 ($\mathcal{M}3$): a Chisquare density χ^2_{10}

 ($\mathcal{M}4$): a Student density t_4

- Some Gaussian mixtures taken from [4] and [11]:

 ($\mathcal{M}5$): $0.5N(-1, 0.3) + 0.5N(1, 0.3)$

 ($\mathcal{M}6$): $0.5N(-2.5, 1) + 0.5N(2.5, 1)$

 ($\mathcal{M}7$): $0.25N(-3, 0.5) + 0.5N(0, 1) + 0.25N(3, 0.5)$

- Gaussian mixtures used in from [9] and taken from [7]

 ($\mathcal{M}8$): the Claw density, $0.5N(0,1) + \sum_{i=0}^{4} \frac{1}{10} N\left(\frac{i}{2} - 1, \frac{1}{10}\right)$

 ($\mathcal{M}9$): the Smooth Comb Density, $\sum_{i=0}^{5} \frac{2^{5-i}}{63} N\left(\frac{65 - 96\frac{1}{2^i}}{21}, \frac{\left(\frac{32}{63}\right)^2}{2^{2i}}\right) =$

 $$\frac{32}{63} N\left(-\frac{31}{21}, \frac{32}{63}\right) + \frac{16}{63} N\left(\frac{17}{21}, \frac{16}{63}\right) + \frac{8}{63} N\left(\frac{41}{21}, \frac{8}{63}\right) + \frac{4}{63} N\left(\frac{53}{21}, \frac{4}{63}\right) + \frac{2}{63} N\left(\frac{59}{21}, \frac{2}{63}\right) + \frac{1}{63} N\left(\frac{62}{21}, \frac{1}{63}\right)$$

- Mixtures density with highly inhomogeneous smoothness as in [9]:

 ($\mathcal{M}10$): $0.5N(0,1) + 0.5 \sum_{i=1}^{10} \mathbf{1}_{\left(\frac{2(i-1)}{T}, \frac{2i-1}{T}\right]}$

 ($\mathcal{M}11$): $0.5N(0,1) + 0.5 \sum_{i=1}^{14} \mathbf{1}_{\left(\frac{2(i-1)}{T}, \frac{2i-1}{T}\right]}$

All the simulations are done with the R software, and for models $\mathcal{M}8$ and $\mathcal{M}9$ we use the benchden package.

We show below in figures 4, 5 and 6, the true densities for the eleven models as well as their estimates obtained using the three algorithms *AggregHist*, *BagHist* and *StackHist* for $n = 500$ observations and $M = 300$ histograms for the two first algorithms.

AggregHist and *BagHist* give more smooth estimators than *StackHist*.

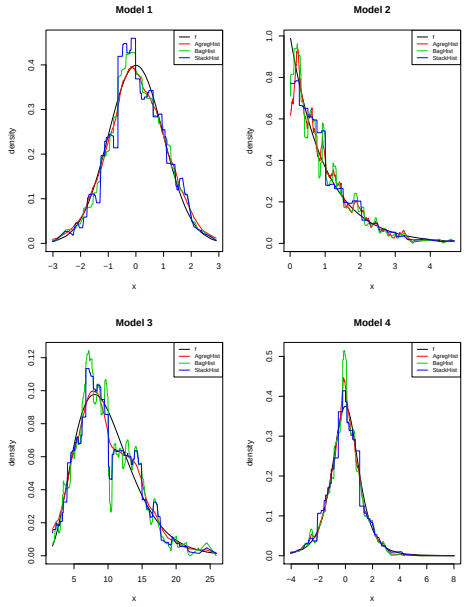


Figure 4: Densities used in simulation models 1 to 4 together with the corresponding histogram and the estimators given by AggregHist, BagHist and StackHist.

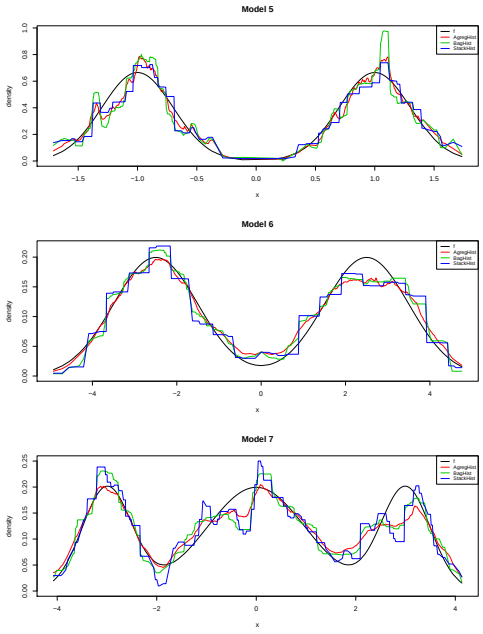


Figure 5: Densities used in simulation models 5 to 7 together with the corresponding histogram and the estimators given by AggregHist, BagHist and StackHist.

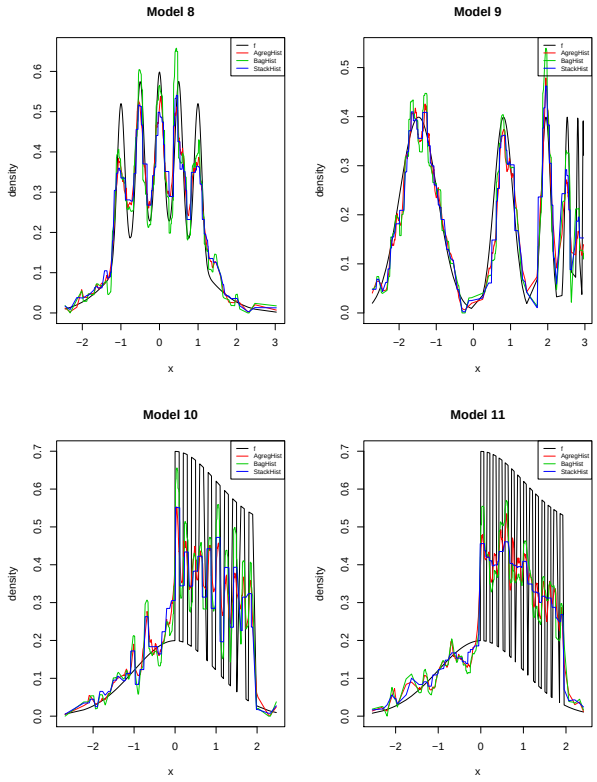


Figure 6: Densities used in simulation models 8 to 11 together with the corresponding histogram and the estimators given by AggregHist, BagHist and StackHist.

4.2 Tuning the algorithms

For the existing algorithms we have used the same values suggested by their corresponding authors :

- For *Stacking*, six kernel density estimators are aggregated, three of them use Gaussian kernels with fixed bandwidths $h = 0.1, 0.2, 0.3$ and three triangular kernels with bandwidths $h = 0.1, 0.2, 0.3$. The number of cross validation samples is fixed to $V = 10$.

- For *AggPure* six kernel density estimators are aggregated having bandwidths 0.001, $0.005, 0.01, 0.05, 0.1$ and 0.5. We use the EM algorithm to optimize the coefficients of the linear combination. The final estimator is a mean over $S = 10$ random splits of the original data set.

- For *Boostkde*, we use 5 steps for the algorithm aggregating kernel density estimators whose bandwidths are optimized using Silverman rule. Normalization of the output

is done using numerical integration. Extensive simulations showed that more steps give rise to less accurate estimators.

Simple one kernel density estimators are also used in our comparisons using optimized bandwidths following Silverman rule (*KdeNrd0*) and unbiased cross validation (*KdeUCV*). For the Histogram, fixed breaks are systematically used and their number is optimized over a fixed grid, retaining the one which maximizes the log likelihood of the obtained histogram over a test sample drawn from the same distribution as the learning sample. The tuning parameters for our algorithms, the number of breakpoints and the value of γ are optimized testing different values for each of them over a fixed grid. We test $10, 20$ and 50 equally spaced breakpoints for each case. For γ we chose the grid $0.5, 1, 1.5, 2, 2.5$. The best combination retained for each model is the one which maximizes the log likelihood over 100 independent test samples drawn from the corresponding model. For $BagHist$ and $AggregHist$ we aggregate $M = 300$ histograms. The optimal values for the histogram and for our algorithms are give in table 1. We denote the optimal number of breaks L_H, L_{BH}, L_{AH} for the Histogram, *BagHist* and *AggregHist* respectively, and γ_{AH} the perturbation coefficient for *AggregHist*.

	$n = 100$				$n = 500$				$n = 1000$			
	L_H	L_{AH}	L_{BH}	γ_{AH}	L_H	L_{AH}	L_{BH}	γ_{AH}	L_H	L_{AH}	L_{BH}	γ_{AH}
$\mathcal{M}1$	50	10	50	1.0	50	10	10	0.5	50	20	20	0.5
$\mathcal{M}2$	50	10	50	0.5	50	50	50	0.5	50	50	50	0.5
$\mathcal{M}3$	50	10	50	0.5	50	10	50	0.5	50	20	50	0.5
$\mathcal{M}4$	50	20	50	0.5	50	50	50	0.5	50	50	50	0.5
$\mathcal{M}5$	50	20	50	1.0	50	50	50	2.0	50	50	50	1.0
$\mathcal{M}6$	50	10	50	1.0	50	20	20	1.0	50	50	20	2.0
$\mathcal{M}7$	50	20	10	0.5	20	20	20	0.5	20	50	20	0.5
$\mathcal{M}8$	50	20	50	0.5	50	50	50	0.5	50	50	50	2.0
$\mathcal{M}9$	50	20	50	0.5	50	50	50	0.5	50	50	50	1.0
$\mathcal{M}10$	50	50	50	0.5	50	50	50	0.5	50	50	50	0.5
$\mathcal{M}11$	50	20	50	0.5	50	50	50	0.5	50	50	50	0.5

Table 1: Optimal parameters values used for our algorithms.

Finally, for *StackHist* we aggregate six histograms having $5, 10, 20, 30, 40$ and 50 equally spaced breakpoints. A ten fold cross validation is used.

4.3 Results

The performance of each model is evaluated using the Mean Integrated Squared Error (MISE). It is estimated as the average of the integrated squared error over 100 simula-

tions. First, for both *AggregHist* and *BagHist* we analyze the effect of the number M of histograms aggregated. Figures 7, 8 and 9 show how the MISE varies when increasing the number of histograms. These graphics show clearly the contribution of the aggregation to the reduction of the MISE. For all the models, the error does not decrease significantly after about 100 iterations.

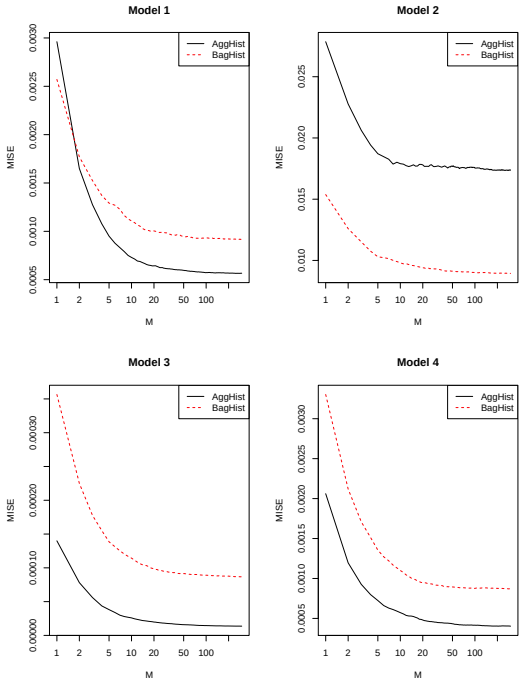


Figure 7: MISE error versus number of aggregated histograms for Models 1 to 4.

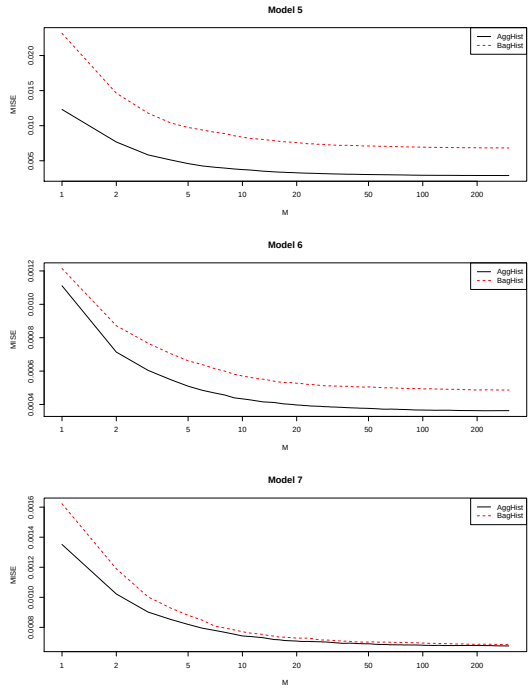


Figure 8: MISE error versus number of aggregated histograms for Models 5 to 7.

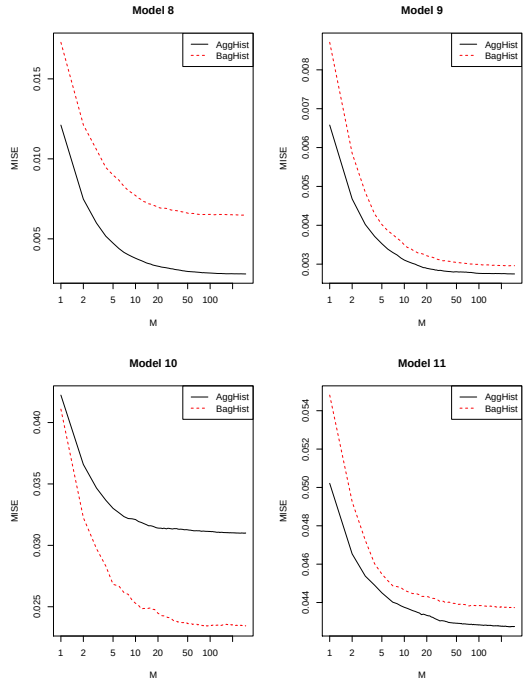


Figure 9: MISE error versus number of aggregated histograms for Models 8 to 11.

We compare now our algorithms *BagHist*, *AggregHist* and *StachHist* to the following methods: the Histogram, *KdeNrd0*, *KdeUCV*, *Stacking*, *AggPure* and *BoostKde*. We have limited the comparisons to the ensemble methods which aggregated non parametric density estimators.

Tables 1 to 3 give the values of $100 \times MISE$ for each method and simulation model for the three values of n. The best performances are indicated in bold.
In most cases, aggregation models have higher accuracy than simple methods like the histogram and KDE. However this is not the case for model $\mathcal{M}3$: KDE has a better performance, that is probably due to border effects. *BoostKde* gives in general better estimates for mixture models. All the methods have better accuracy in general when increasing the sample size n. *BagHist* and *AggregHist* always outperform the optimal Histogram and in general both have better accuracy than *StackHist*. For $n = 100$ *BagHist* and *AggregHist* outperform the other algorithms over the complicated models ($\mathcal{M}8 - \mathcal{M}10$). For $n = 500$ and $n = 1000$, *AggPure* outperforms the other methods for the last two models. Although our algorithms do not always outperform the other methods, their precision are not far from the best ones.

Table 2: $n = 100, M = 300, 100\times$ MISE

100	Histogram	KdeNrd0	KdeUCV	BoostKde	AggPure	Stacking	StackHist	BagHist	AggregHist
$\mathcal{M}1$	2.870	**0.193**	0.221	0.831	0.365	0.233	0.511	2.250	0.209
$\mathcal{M}2$	4.820	7.040	4.710	9.320	2.750	2.610	1.580	4.330	**1.560**
$\mathcal{M}3$	0.162	**0.012**	0.013	0.035	0.106	0.061	0.032	0.127	0.015
$\mathcal{M}4$	1.380	**0.181**	0.219	1.460	0.318	0.202	0.370	1.130	0.220
$\mathcal{M}5$	6.700	6.410	3.070	**0.387**	1.040	0.856	1.940	5.040	0.911
$\mathcal{M}6$	0.739	0.149	0.082	**0.056**	0.189	0.111	0.178	0.550	0.080
$\mathcal{M}7$	0.820	0.278	0.156	**0.107**	0.163	0.110	0.168	0.126	0.128
$\mathcal{M}8$	4.310	2.280	2.130	3.840	1.910	1.620	1.850	3.390	**1.560**
$\mathcal{M}9$	2.460	2.330	1.350	1.350	0.967	0.927	1.130	1.950	**0.832**
$\mathcal{M}10$	5.340	6.750	6.940	6.090	4.480	4.700	4.780	**3.640**	3.720
$\mathcal{M}11$	6.130	6.620	6.840	5.590	5.160	5.650	5.590	**4.940**	5.130

Table 3: $n = 500, M = 300, 100\times$ MISE

500	Histogram	KdeNrd0	KdeUCV	BoostKde	AggPure	Stacking	StackHist	BagHist	AggregHist
$\mathcal{M}1$	0.474	0.085	0.100	0.098	0.076	**0.054**	0.154	0.085	0.056
$\mathcal{M}2$	0.786	5.980	2.840	6.930	1.280	1.120	**0.545**	0.701	0.720
$\mathcal{M}3$	0.024	**0.003**	**0.003**	0.029	0.020	0.012	0.009	0.018	0.004
$\mathcal{M}4$	0.180	0.099	0.113	1.160	0.059	**0.043**	0.114	0.131	0.072
$\mathcal{M}5$	1.190	5.070	2.090	**0.132**	0.430	0.254	0.591	0.870	0.424
$\mathcal{M}6$	0.133	0.094	0.037	**0.014**	0.036	0.024	0.056	0.040	0.024
$\mathcal{M}7$	0.079	0.206	0.085	**0.021**	0.036	0.035	0.067	0.050	0.034
$\mathcal{M}8$	0.859	2.120	1.770	2.010	0.768	0.588	0.683	0.547	**0.531**
$\mathcal{M}9$	0.594	2.020	0.896	1.070	0.443	0.393	0.428	0.488	**0.377**
$\mathcal{M}10$	3.680	6.380	5.690	5.640	**2.260**	2.740	3.790	2.800	3.420
$\mathcal{M}11$	4.700	6.350	6.120	5.880	**2.760**	4.430	4.900	4.160	4.000

Table 4: $n = 1000, M = 300, 100\times$ MISE

1000	Histogram	KdeNrd0	KdeUCV	BoostKde	AggPure	Stacking	StackHist	BagHist	AggregHist
$\mathcal{M}1$	0.207	0.066	0.081	0.136	0.039	**0.032**	0.087	0.064	0.036
$\mathcal{M}2$	0.374	5.690	2.500	6.090	0.896	0.850	0.365	**0.339**	0.676
$\mathcal{M}3$	0.011	**0.002**	**0.002**	0.014	0.009	0.006	0.005	0.008	**0.002**
$\mathcal{M}4$	0.099	0.071	0.076	0.816	0.036	**0.028**	0.074	0.062	0.035
$\mathcal{M}5$	0.593	4.530	1.730	**0.062**	0.258	0.146	0.318	0.420	0.216
$\mathcal{M}6$	0.064	0.078	0.030	**0.008**	0.018	0.013	0.032	0.023	0.024
$\mathcal{M}7$	0.054	0.181	0.066	**0.012**	0.021	0.020	0.042	0.031	0.031
$\mathcal{M}8$	0.590	2.070	1.690	1.530	0.498	0.381	0.506	**0.336**	0.506
$\mathcal{M}9$	0.406	1.890	0.767	0.902	0.307	**0.286**	0.331	0.332	0.342
$\mathcal{M}10$	3.830	6.220	5.170	5.360	**1.600**	2.450	3.880	2.980	3.490
$\mathcal{M}11$	4.700	6.220	5.590	5.340	**1.950**	4.100	4.790	4.380	4.070

5 Conclusion

In this work we present three new algorithms for density estimation aggregating histograms. Two of them aggregate histograms over bootstrap samples of the data or randomly perturbed breakpoints. The third is a simple adaptation of the stacking algorithm where histograms are used instead of kernel density estimators. We have shown using extensive simulations that these algorithms and the other ensembles techniques have better accuracy than the histogram or KDE. The first two algorithms *BagHist* and *AggregHist*

are very simple to implement, depend on very few parameters, and their computation complexity is proportional to that of a histogram. Theoretical properties of these algorithms are under study. Most of the algorithms described in this work may be easily generalized to the multivariate case.

References

[1] L. Breiman. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

[2] L. Breiman. Stacked regression. *Machine Learning*, 24(1):49–64, 1996.

[3] L. Breiman. Random forests. *Machine Learning*, 45(1):5–32, October 2001.

[4] M. Di Marzio and C.C. Taylor. Boosting kernel density estimates: A bias reduction technique? *Biometrika*, 91(1):226–233, 2004.

[5] Y. Freund. Boosting a weak learning algorithm by majority. *Information and Computation*, 121,2:256–285, 1995.

[6] M. Jones, O. Linton, and J. Nielsen. A simple bias reduction method for density estimation. *Biometrika*, 82:327–338, 1995.

[7] J.S Marron and M.P. Wand. Exact mean integrated square error. *The Annals of Statistics*, 20(2):712–736, 1992.

[8] G. Ridgeway. Looking for lumps: boosting and bagging for density estimation. *Comput. Stat. Data Anal.*, 38(4):379–392, 2002.

[9] P. Rigollet and A. B. Tsybakov. Linear and convex aggregation of density estimators. *Math. Methods Statist.*, 16(3):260–280, 2007.

[10] S. Rosset and E. Segal. Boosting density estimation. In *In Advances in Neural Information Processing Systems 15*, pages 641–648. MIT Press, 2002.

[11] P. Smyth and D. Wolpert. Linearly combining density estimators via stacking. *Mach. Learn.*, 36(1-2):59–83, 1999.

[12] X. Song, K. Yang, and M. Pavel. Density boosting for gaussian mixtures. *Neural Information Processing*, 3316:508–515, 2004.

[13] D.H. Wolpert. Stacked Generalization. *Neural Networks*, 5:241–259, 1992.